②

AD-A222 430

A WORKBENCH FOR DISCOVERING TASK-SPECIFIC THEORIES OF LEARNING

Technical Report PCG - 18

Kurt VanLehn
Departments of Psychology & Computer Science

DEPARTMENT
of
PSYCHOLOGY

DTIC
S ELECTE
JUN 0 6 1990
E D

Carnegie Mellon University

90 06 05 056

2

A WORKBENCH FOR DISCOVERING
# TASK-SPECIFIC THEORIES OF LEARNING

Technical Report PCG - 18

Kurt   VanLehn
Departments of Psychology & Computer Science

Carnegie Mellon University
Pittsburgh, PA 15213 U.S.A.

3 March 1989

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION	1b. RESTRICTIVE MARKINGS
Unclassified	

2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION / AVAILABILITY OF REPORT
	Approved for public release;
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE	Distribution unlimited

4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S)
PCG - 18	Same as Performing Organization

6a. NAME OF PERFORMING ORGANIZATION	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION
Carnegie Mellon University		Personnel and Training Research Office of Naval Research (Code 1142PT)

6c. ADDRESS (City, State, and ZIP Code)	7b. ADDRESS (City, State, and ZIP Code)
Department of Psychology Pittsburgh, Pennsylvania 15213	800 N. Quincy Street Arlington, VA 22217-5000

8a. NAME OF FUNDING / SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
Same as Monitoring Organization		N00014-88-K-0086

8c. ADDRESS (City, State, and ZIP Code)	10 SOURCE OF FUNDING NUMBERS			
	PROGRAM ELEMENT NO	PROJECT NO	TASK NO.	WORK UNIT ACCESSION NO
	N/A	N/A	N/A	N/A

11 TITLE (Include Security Classification)

A Workbench for Discovering Task-Specific Theories of Learning

12. PERSONAL AUTHOR(S) VanLehn, Kurt Alan

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Technical	FROM 88Jan01 TO 90Dec31	1989 March 03	14

16 SUPPLEMENTARY NOTATION In T.O'Shea & E. Scanlon (Eds.) "Proceedings of the NATO Workshop on Advanced Educational Technology"

17	COSATI CODES		18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	Cognitive Task analysis, instructional design, cognitive modelling, learning , (⊆⎹ ⁻ ⌐)

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

Very little of our cognitive behavior (as opposed to more peripheral behaviors) is determined by the fixed, unchangeable parts of our mind. Cognitive behaviors seem to be determined by our knowledge and the environment itself. Thus, recent work in AI on developing models of the fixed part of the mind (the cognitive architecture) will not be of much use to educators who wish to perform a cognitive task analysis of their subject matter before designing instruction for it. However, it seems technically feasible, given the state of the art in Machine Learning, to build semi-automated data analysis packages that can be added to a cognitive architecture, thus creating a "workbench" for performing cognitive task analysis. Such tools becomes particularly interesting when used to analyze not only the experts' understanding of the subject matter, but to trace the development of a student's understanding. In this application, the workbench can lead to task-specific theories of learning which are much more informative to educators than most learning theories have been.

20. DISTRIBUTION / AVAILABILITY OF ABSTRACT	21 ABSTRACT SECURITY CLASSIFICATION
☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ DTIC USERS	

22a NAME OF RESPONSIBLE INDIVIDUAL	22b TELEPHONE (Include Area Code)	22c. OFFICE SYMBOL
Susan Chipman	(202) 696-4322	1142 PT

1. Introduction

One long-standing aspiration of cognitive science is that education would benefit from the building of learning theories that are expressed, at least partially, as Artificial Intelligence (AI) programs. I have built several such programs (VanLehn, 1987; VanLehn & Ball, 19??), and others have built many more (Anderson, 1983; Newell, 19??; Holland, Holyoak, Nisbett & Thagard, 1986; Anzai, 1987; Ohlsson, 1987). Although such work has profoundly changed our image of competence and intelligence, and that change has begun to seep into the educational system, it is fairly clear now that the resulting programs/theories have not had as much *direct* effect on education and training as could be desired. This paper examines the reasons why and suggests a new research direction based on that analysis.

The basic problem is that there seems to be an unavoidable tradeoff between the generality of learning theories and their utility to educators. Let us examine this tradeoff by starting with some recent general theories of learning and seeing what utility they have for education.

SOAR (Newell, 19??; Laird, Newell, & Rosenbloom, 1987) and ACT* (Anderson, 1983; Anderson, 1987) aim to be universal theories of cognition. Their goal is to describe only the aspects of skill acquisition that are common to the acquisition of *all* skills. These theories are well suited for some purposes. Some examples are:

- explanations of speed and error patterns in transcription typing (John, 1988),

- explanations of the power-law increase in speed and accuracy that invariably accompany extensive practice (Rosenbloom & Newell, 1987; Anderson, 1983),

- explanations of transfer, as measured by savings in learning time caused by prior training on a similar skill (Singley & Anderson, 1985; Singley & Anderson, 19??; Kessler, 1988).

However, the mechanisms of ACT* and SOAR do not in themselves tell us much about the students' *initial* acquisition of the skill. For instance, they do not tell us how students will read an instructi the effects of examples, nor the impact of sp⟨ ⟩ pre-existing conceptual knowledge, nor the importance of having mental models in task domains ⟨ ⟩ admit them, and so forth.

This is not an oversight on the part of the authors of ACT* and SOAR, but arises from the fact that initial acquisition of a skill seems to be a form of problem solving. Students, while engaged in various pedagogical activities such as studying a text or working some exercises, occasionally discover that their knowledge is incomplete or mistaken. This is a problem. They know many methods for solving the problem of ignorance, and different students may know different methods.[1]

As always in problem solving, the behavior of the subjects is determined mostly by the nature of the problem and the particulars of their knowledge. Neither of these is specified by ACT* or SOAR, as they aim to describe only the *universal* aspects of cognition. However, ACT* and SOAR should be consistent with the observed behavior in that one should be able to specify (as ACT* or SOAR programs) a model of the individual subjects' knowledge and the task environment that will cause the architectures to accurately simulate his or her behavior. Presumably, the particulars of ACT* and SOAR put some constraints on the specification of the knowledge, but the constraints imposed by the nature of the task are much stronger.

To put it differently, suppose an educator who is interested in teaching thermodynamics is not sure which of several ways of learning is typically used by thermodynamics students or could potentially be used by them. Trying these various options out on ACT* and SOAR will not reduce the educator's uncertainty one bit, because the architectures will probably be consistent with all learning methods the educator is likely to consider. In short, because these architectures aim at universality, they turn out to pretty useless as constraints on task-specific theories of initial skill acquisition.

To put the same point a third way, one view of pedagogy (Anderson et al., 1984) is that a sufficient teaching method (but not, of course a necessary one) is to:

1. formalize as production rules (or some other type of rule) exactly what the students need to know in order to perform competently, and

2. design a curriculum whose lessons introduces these rules in small batches (cf. VanLehn, 1983), and

3. design lessons that explain the rules clearly and provide sufficient practice on applying them. (Immediate feedback is seen as particularly important for catching misunderstandings and rectifying them, but it is not essential to this method.)

The critical step in this teaching method is the task analysis that takes place in the first step. Task analysis is driven almost exclusively by the subject matter of the task domain. General cognitive theories, such as ACT*, provide a notation for the rules, but otherwise offer little guidance to the person conducting the analysis.

2. The essential problem, and three possible solutions

These deficiencies are not a fault of ACT* and SOAR per se. Rather, it seems that very little of our cognitive behavior (as opposed to more peripheral behaviors) is determined by the fixed, unchangeable parts of our mind. Cognitive behaviors seem to be determined by our knowledge and the environment itself. Moreover, knowledge acquisition is a cognitive behavior, which is itself determined mostly by knowledge and the environment. To put it in more traditional terms, because we humans are a highly adaptive species (i.e., we mold our behavior to fit the environment), our higher level behavior is determined mostly by our history of interaction with the environment (our knowledge) and by the environment at hand.

Unpacking the recursion here, it seems that the ultimate determinant of cognitive behaviors is the person's environment. (This is, of course, a gross simplification -- I am not proposing a tabla rasa here.) Presumably, one could explain cognitive behavior by omitting descriptions of the various cycles of knowledge acquisition, etc. and just examine the relationship between the environment and cognitive behavior.[2] Although this is one logically possible way to predict human behavior, I suspect that such an explanation would be cumbersome and inaccurate, so I would not recommend pursuing it.

Logically, the only other option is to incorporate the environment into the theory. Thus, for example, a theory of physics learning would include task-specific terms like "forces" and "equations." Such theories blend psychology and the particulars of a task domain. In order to illustrate the notion of task-specific theories, let us examine some simple ones. The task of arithmetic calculation is fairly well understood. It divides cleanly into recall of arithmetic facts, such as 17–9=8, and execution of arithmetic algorithms, such as the algorithm for subtracting two multidigit numbers. We will consider a task-specific theory for recall and a task-specific theory for execution.

Siegler (Siegler & Shrager, 1984; Siegler, 19??a; Siegler, 19??b) has developed specific models of how students "recall" arithmetic facts. Each model has parameters that can be fit to a individual subject's behavior, thus providing both a test of the models and a way to forecast the subject's behavior. Each model is specific to one type of arithmetic operation, but they are all consistent with his general theory of strategy selection, which features a specific procedure for trading off retrieval and reconstruction of the item to be recalled. Reconstruction, in this context, might consist of using counting to generate an addition fact. Moreover, the general theory specifies how memory traces are strengthened by practice, thus leading to the dominance of memory retrieval over reconstruction that

characterizes the competent student's performance. Siegler's theory of recall seems quite general, for it has been successfully applied to analyze acquisition of spelling rules (Siegler, personal communication) as well as the major arithmetic operations. Of course, it is not as general as ACT* or SOAR, but it serves nicely as a simple illustration of the difference between a general theory, a task-specific theory/model (e.g., the model for addition, which has explicit reconstruction strategies for arithmetic facts), and a subject-specific model (the addition model, with its parameters fit to a given subject's data). Siegler's task-specific models are specific enough that one can envision designing a curriculum around them, and Siegler has recently begun to do just that (Siegler, personal communication).

My colleagues and I have developed models of the algorithms for multidigit arithmetic, concentrating especially on subtraction (Brown & VanLehn, 1980; VanLehn, 1983b; VanLehn, 19??; VanLehn, Ball & Kowalski, 1988). There is a general theory, which distinguishes between normal execution of a procedure and "error handling." According to the theory, when people reach an impasse, perhaps because their knowledge of the procedure is incomplete and they can not decide what to do next, they treat the impasse itself as a problem and attempt to resolve it. One impasse-resolving strategy is to ask for help or to consult a textbook. Another is to search through one's earlier work looking for an inadvertent error. These strategies depend strongly on the particulars of situation that the students are in and on their knowledge of the task domain. Another hypothesis of the general theory is that learning occurs whenever the resolution of an impasse is summarized and stored in memory as a new rule (VanLehn, 1988a). The general theory has been tested by developing a task-specific theory/model of subtraction (VanLehn, 1983b; VanLehn, 19??). The model has been fit to individual subjects' error data. The task-specific model makes predictions about pedagogies for subtraction, some of which have been tested (VanLehn, 1988b). This work again illustrates the difference between a general theory, which offers little specific guidance to educators, and task-specific theories/models, which provide crisp suggestions.

Neither of the "general" theories just mentioned are as general as ACT* or SOAR, so a better view of the world is to see theories as arranged in some kind of generalization hierarchy. SOAR, for instance, is a straightforward generalization of both Siegler's theory and mine, because it generalizes the notion of an "impasse" to cover both failures due to memory retrieval and failures due to flawed knowledge. On the other hand, SOAR offers even less guidance to educators than either Siegler's theory or mine, just because it has more generality. So the same generality-power tradeoff is evident, even though the

binary distinction between general theories and task-specific ones has dissolved into a generalization hierarchy. Although I will continue to speak of "general" versus "task-specific" theories, one should keep in mind that this is a simplification.

It seems that task-specific theories offer a viable option for guiding pedagogy. But unfortunately, task-specific theories offer little help to people who are interested in other tasks (or at least, that is how the theories are treated: theories of arithmetic are pretty much ignored by everyone except those interested in arithmetic). Thus, while task-specific theories are much more h 'pful to some educators than general theories, they are not helpful to very many educators.

This leads to a third option (the first two were environmental theories and task-specific theories), which is to formulate a *method* for generating task-specific theories. Traditionally, a method is a prescription of the kinds of experiments to run, the kinds of analyses to make and the kinds of conclusions to draw. The later two items are actually a weak task-general theory. It is weak because it does not foreordain the conclusions, but merely provides some ideas or even some notations for stating the task-specific theory. To put it differently, a method provides (1) a general theory and (2) a means of instantiating the theory to fit a task domain, thus formulating a task-specific theory.

There are methods in education, but I believe it is fair to say that all of them are oriented towards prescribing instruction rather than constructing learning theories. The social sciences contain many descriptive methods, such as factor analysis and its associated theory of intelligence, or structural linguistics and its associated theory of syntax. However, as far as I know, there is no method for formulating task-specific theories of learning.

This does not bode well for a project aimed at formulating such a method. All the arguments presented above depend only on ancient concepts, such as the distinction between knowledge and its application. These arguments lead more or less inevitably to the project of formulating a method. Surely someone in the long history of education and psychology must have tried to formulate such a method. Maybe they tried and failed. Maybe such a method is just not feasible.

Some recent results in AI indicate that a method for formulating task-specific theories may indeed be feasible. Most of the work is aimed at replicating the reasoning processes behind human scientific discovery (Langley, Simon, Bradshaw & Zytkow, 1987; Kulkarni & Simon, 1988; Shavlik, 1986).

Although there is no denying that these programs produce the same hypotheses and experimental demonstrations that the human scientists did, there are still grave doubts about whether the simplifications assumed by these models are too strong. Pessimists would say that the machine discovery programs are not particularly intelligent, but the people who chose the simplifications for them were very intelligent. Since the pessimists could turn out to be right, it is prudent for those who wish to apply this new machine discovery technology to assume that a practical machine discovery system has a scientist/user who selects the simplifications and oversees the machine's reasoning. To put it crudely again, although the machine discovery work may or may not be able to build a mechanical scientist, it probably can build a mechanical research assistant. Such a tool could play a key role in a method for formulating task-specific theories of learning.

In short, it seems that the most promising option for finding theories of learning that are really useful to educators is to formulate a method that combines the talents of people and machine discovery programs in order generate task-specific theories of learning. This is a research option that I think should be pursued.

3. Workbenches: existing and proposed

Calling the research product a "method" makes it sound like a step-by-step prescription of how to construct a theory. I do not think *that* kind of method is feasible. What I have in mind is a set of integrated computer-based tools for analyzing data and building models. Such a "scientist's workbench" would be based on some task-general theory, such as ACT* or SOAR, or perhaps some moderately general theory, like Siegler's or mine. This section discusses some examples.

CIRRUS (VanLehn & Garlick, 1987; Kowalski & VanLehn, 1988) is a workbench based on my theory about how people execute cognitive. In addition to the hypotheses mentioned above, the theory includes the hypotheses that people are free to pick any goal that they can recall as the next goal to attend to, and their knowledge includes some policies concerning what types of goals to attend to in what situations (VanLehn, Ball & Kowalski, 1988).[3] CIRRUS is designed for analyzing protocol data within the framework of the theory by building a runable simulation and comparing its behavior to the given protocol.[4] Students' policies about goal selection are formalized as a set of goal selection preferences of the form "If condition C holds, then prefer goals of type A over goals of type B." The simulator uses such preferences to sort a list of pending goals and choose the goal that is preferred above all others. To use

CIRRUS, the theorist must input a procedure, written in the knowledge representation language of the theory, that lacks goal selection preferences. CIRRUS must also be given primitives from which goal selection preferences can be built. Given a protocol, CIRRUS builds goal selection preferences that allow a maximally accurate simulation of the data. To put it in more traditional terms, CIRRUS takes a model with one parameter, and fits it to the given data. However, both the model and the parameter are non-numeric.

When my collaborators and I use CIRRUS, we find it necessary to refine the model given to it many times before we are finally happy with the analysis it yields. Typically, we analyze one subject's data in some detail, then start our analysis of the next subject using the model developed for the first subject. After several subjects have been analyzed, commonalties in the subject-specific models emerge. At that point, we build a subject-general model and install parameters (typically, a system of switches that turn rules off and on) in order to capture the between-subjects variation. We stop the analysis when all the subjects have been analyzed and one subject-general model has been found. One of the model's parameters, the set of goal-selection procedures, is fit automatically by CIRRUS; the other parameters, which were created during the model refinement process, are fit by hand. This refinement process can be viewed as finding a theory that is specific to the task under analysis but general across subjects. In this fashion, CIRRUS helps the scientist/user discover a task-specific theory/model.

ACM (Ohlsson & Langley, 1985; Ohlsson & Langley, 1988) is similar to CIRRUS. It is based on the theory that problem solving is search through a problem space. It takes as its model a specific problem space, and builds a set of operator selection heuristics that will cause search through this problem space to simulate answer data given to the program.

SAPA (Bhaskar & Simon, 1977) is somewhat like ACM, in that it is based on the theory of problem solving as search through a problem space. However, it does not to actually build a set of search heuristics that fit some data given to it. It already has some search heuristics in it, along with a particular problem space.[5] These search heuristics are intended, I suppose, to represent those of a prototypical subject's. At each cycle of the search, SAPA asks the user if the inference it has just made corresponds to the protocol. If it does, then the built-in, fully parameterized model is upheld. If not, then SAPA checks to see if the parameterization is wrong -- i.e., it has the right problem space but the wrong heuristics for that subject. It performs this check by suggesting alternatives until the user indicates that it has found one that corresponds to the protocol. If none of SAPA's suggestions work, then the problem

space is deemed faulty, because no parameterization of the model will fit the data. Bhashkar and Simon used SAPA to test their task-specific theory of thermodynamics problem solving, and to test their model of a prototypical student's search heuristics.

All these workbenches, as well as several others (e.g., DEBUGGY (Burton, 1982), TETRAD (Glymour, Scheines, Spirtes & Kelly, 1987) and METADENDRAL (Lindsay, Buchanan, Feigenbaum & Lederberg, 1980)) have three components: (1) a general theory that is so deeply embedded in the workbench that it can not be changed, (2) a underdetermined model given to the workbench by the user, such as a problem space for thermodynamics problem solving, and (3) a process that fits the model to the data, making it more deterministic. The theorist tinkers with the underdetermined model in order to get a fitted model that analyzes the data satisfactorily. The result is a model that is both a generalization over several (hopefully, many) subjects' data and a specialization of the general theory. The model can be considered a task-specific theory.

Of course, such a model is interesting only to the extent that that task is interesting. Educators are interested in learning, but CIRRUS, ACM and SAPA all assume that learning does not occur during the protocols they are analyzing. Thus, they could be used in a longitudinal study to model snapshots of the learner's development, but they can not model the learning process itself. This leads to a proposal to build a workbench that can model the learning process.

I am currently involved in building a scaled-up version of CIRRUS, called CASCADE. CASCADE is being built in order to analyze a very large data set, donated by Micki Chi (Chi, Bassok, Lewis, Reimann & Glaser, 19??). The data consist of 8 protocols, each about 200 pages long. They were collected from students studying the first four chapters of a college physics textbook. The protocols record the learning that a typical college student would undergo in the first few weeks of a college physics course.

4. Expected benefits of the proposed research

The most important application of the proposed technology is providing a "front end" to projects that create training systems. According to Anderson, the first step in developing a training system is to analyze the task domain to see what good students should know when they have completed their training (Anderson et al., 1984). Workbenches such as CASCADE are intended to help a designer perform such a task analysis. Although this section suggests a few other benefits that might accrue, one should keep in mind that the main benefit is technological assistance in task analysis.

The task-specific, subject-general model that is created on the workbench could be the starting point of the development of a student modeler for an intelligent tutoring systems. Also, the data analysis tools developed as parts of the workbench could be used as parts of the diagnostic module of an intelligent tutoring system.

The mere process of analyzing students' learning in the face of the given instructional material will usually reveal defects in the material that can be easily remedied. Anderson, for instance, has a written a textbook on Lisp based on his task analysis. Since the analysis had only gotten as far as recursion when the book was written, the last five chapters in the text were not based on a task analysis. Anderson comments: "Since the writing of the book we have slowly began to create tutor material corresponding to those chapters. As we have done so we have started to realize the inadequacy of the information in the last five chapters." (Anderson, in press, chapter 4). It is significant that task analysis of the initial segment of the curriculum, even by someone like Anderson, was not sufficient preparation for writing an adequate material for the second segment. It seems that there is no substitute for formal task analysis, even if the intended training vehicle is "just" as textbook.

Once a task-specific model of the student has been constructed, it often suggests new pedagogical strategies. Given the model, some will seem clearly beneficial. However, pedagogies whose benefits are less certain can be simulated; if the model is psychologically accurate, and the proposed benefit helps the model learn, then human students should learn better as well. For instance, on the basis of Siegler's model of addition, it seems that under certain circumstances, supervised drill can take advantage of the commutativity of addition and only teach half the addition facts. Unsupervised drill on the other half should suffice for learning them. This pedagogical regime should be tested on his model before being tried in the classroom.

So far, the importance of this work to education has been stressed. But there are other potential beneficiaries as well. Machine learning has recently turned towards scientific discovery as a source of new problems. Because a workbench is a program that participates in scientific discovery, it should be of some interest to research on discovery. One can even imagine taking protocols of scientists while they use it in order to understand the discovery-making process better.

In protocols of students involved in learning new material, such as the ones being analyzed by the CASCADE project, there are many instances of students making discoveries. These discoveries might suggest discovery methods that could be developed into full-fledged machine learning techniques.

Looking further ahead, machine learning has not yet produced interactive learners that can hold up their end of a training dialog with their trainer. Formal work in the Valiant framework ("PAC learning") indicates that such interactivity is necessary for tractable learning (Valiant, 1984), so eventually machine learning will have to build such interactive learners if it is to live up to its promises of delivering systems that acquire knowledge for expert systems. The current protocol studies show how interaction proceeds with human students. That should suggest styles of interaction to machine learning researchers.

Turning now to the benefits for psychology, we start with the traditional observation that applications usually push theories towards completion because application efforts do not have the luxury of ignoring parts of human behavior that are difficult to explain. This application of cognitive theory will certainly push it towards completion. For instance, the physics task domain is richer in conceptual material than other task domains, such as Lisp and geometry, that have been studied. Thus, the development of a task-specific theory in physics should illuminate the interaction between conceptual and procedural learning.

I have concentrated on workbenches for analyzing protocol data because such data will push cognitive theory along by explicating the mapping between theoretical events, such as impasses, and visible types of human behavior. There are few published comparisons of protocols and models as detailed as the analyses in Human Problem Solving (Newell & Simon, 1972), and none that compare models and students who are learning. The CASCADE project, and others like it, should yield the first fine-grained analysis of human learning. From such analyses, we ought to uncover some unexpected theoretical problems, as well as strengthen known weak spots in the theory.

References

Anderson, J. R. (1983). *The Architecture of Cognition*. Cambridge, MA: Harvard.

Anderson, J.R. (1987). Skill acquisition: Compilation of weak-method problem soutions. *Psychological Review, 94*(2), 192-210.

Anderson, J.R. (in press, 19??). *The Adaptive Character of Thought*. Cambridge, MA: MIT Press.

Anderson, J. R., Boyle, C. F., Farrell, R. & Reiser, B. J. (1984). Cognitive principles in the design of computer tutors. In *Proceedings of the Cognitive Science Society Conference*. Hillsdale, NJ: Erlbaum.

Anzai, Y. (1987). Doing, understanding and learning in problem solving. In D.Klahr, P.Langley & R.Neches (Ed.), *Production System Models of Learning and Development*. Cambridge, MA: MIT Press.

Bhaskar, R. & Simon, H. A. (1977). Problem solving in a semantically rich domains: An example from engineering thermodynamics. *Cognitive Science, 1*, 193-215.

Brown, J. S. & VanLehn, K. (1980). Repair Theory: A generative theory of bugs in procedural skills. *Cognitive Science, 4*, 379-426.

Burton, R. B. (1982). Diagnosing bugs in a simple procedural skill. In D. H. Sleeman & J. S. Brown (Eds.), *Intelligent Tutoring Systems*. New York: Academic. 157-183.

Chi, M.T.H., Bassok, M., Lewis, M., Reimann, P. & Glaser, R. (in press, 19??). Learning problem solving skills from studying examples. *Cognitive Science*, .

Glymour, C., Scheines, R., Spirtes, P. & Kelly, K. (1987). *Discovering Causal Structure*. Orlando, FL: Academic Press.

Holland, J.H., Holyoak, K.J., Nisbett, R.E. & Thagard, P.R. (1986). *Induction: Process of Inference, Learning and Discovery*. Cambridge, MA: MIT Press.

John, B. (1988). *Contributions to engineering models of human-computer interaction*. Doctoral dissertation, Dept. of Psychology, Carnegie-Mellon University,

Kessler, C.M. (1988). *Transfer of Programming Skills in Novice Lisp Learners*. Doctoral dissertation, Dept. of Psychology, Carnegie-Mellon University,

Kowalski, B. & VanLehn, K. (1988). *Inducing subject models from protocol data* (Tech. Rep. PCG-15). Dept. of Psychology, Carnegie-Mellon University.

Kulkarni, D. & Simon, H.A. (1988). The process of scientific discovery: The strategy of experimentation. *Cognitive Science, 12*, 139-175.

Laird, J. E., Newell, A., and Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, *33*, 1-64.

Langley, P., Simon, H.A., Bradshaw, G.L., & Zytkow, J.M. (1987). *Scientific Discovery: Computational Explorations of the Creative Process*. Cambridge, MA: MIT Press.

Lindsay, R., Buchanan, B.G., Feigenbaum, E.A. & Lederberg, J. (1980). *DENDRAL*. New York: McGraw-Hill.

Newell, A. (in prep., 19??). *Universal Theories of Cognition*. ?addr?: ?pub.

Newell, A. & Simon, H. A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

Ohlsson, S. (1987). Truth versus appropriateness: Relating declarative to procedural knowledge. In D.Klahr, P.Langley & R. Neches (Ed.), *Production System Models of Learning and Development*. Cambridge, MA: MIT Press.

Ohlsson, S. & Langley, P. (1985). *Identifying solution paths in cognitive diagnosis* (Technical Report CMU-RI-TR-84-7). Pittsburgh, PA: Robotic Institute, CMU.

Ohlsson, S. & Langley, P. (1988). Psychological evaluation of path hypotheses in cognitive diagnosis. In H. Mandl & A. Lesgold (Ed.), *Learning Issues for Intelligent Tutoring Systems*. New York, NY: Springer Verlag.

Rosenbloom, P. & Newell, A. (1987). Learning by chunking: A production system model of practice. In D.Klahr, P.Langley & R. Neches (Ed.), *Production System Models of Learning and Development*. Cambridge, MA: MIT Press.

Shavlik, J.W. (1986). Learning classical physics. In T.M. Mitchell, J.G. Carbonell & R.S. Michalski (Ed.), *Machine Learning: A Guide to Current Research*. Hingham, MA: Kluwer.

Siegler, R.S. (in press, 19??). Strategy choices in subtraction. In J.Sloboda & D.Rogers (Ed.), *Cognitive Processes in Mathematics*. Oxford, UK: Oxford University Press.

Siegler, R.S. (in press, 19??). Strategy choice procedures and the development of multiplication skill. *Journal of Experimental Psychology: General*, .

Siegler, R. S. & Shrager, J. (1984). Strategy choices in addition: How do children know what to do? In C. Sophian (Ed.), *Origins of Cognitive Skill*. Hillsdale, NJ: Erlbaum.

Singley, M.K. & Anderson J.R. (1985). The transfer of text-editing skill. *International Journal of Man-Machine Studies*, *22*, 403-423.

Singley, M.K. & Anderson, J.R. (in press, 19??). A keystroke analysis of learning and transfer in text editing. *Human-Computer Interaction*.

Valiant, L.G. (1984). A theory of the learnable. *Communications of the ACM, 27*(11), 1134-1142.

VanLehn, K. (1983). *Felicity conditions for human skill acquisition: Validating an AI-based theory* (Tech. Report CIS-21). Xerox Palo Alto Research Center.

VanLehn, K. (1983). Human skill acquisition: Theory, model and psychological validation. In *Proceedings of AAAI-83.* Los Altos, CA: Morgan Kaufmann,

VanLehn, K. (1987). Learning one subprocedure per lesson. *Artificial Intelligence, 31*(1), 1-40.

VanLehn, K. (1988). Toward a theory of impasse-driven learning. In H. Mandl & A. Lesgold (Ed.), *Learning Issues for Intelligent Tutoring Systems.* New York, NY: Springer Verlag.

VanLehn, K. (1988). *Tutoring beats felicity conditions for cognitive skill acquisition* (Tech. Rep. PCG-17). Department of Psychology, Carnegie-Mellon University. in preparation.

VanLehn, K. (in press, 19??). *Mind Bugs: The origins of procedural misconceptions.* Cambridge, MA: MIT Press.

VanLehn, K. & Ball, W. (in press, 19??). Teton: A large-grained architecture for studying learning. In VanLehn, K. (Ed.), *Architectures for Intelligence.* Hillsdale, NJ: Erlbaum.

VanLehn, K. & Garlick, S. (1987). Cirrus: an automated protocol analysis tool. In Langley, P. (Ed.), *Proceedings of the Fourth Machine Learning Workshop.* Los Altos, CA: Morgan-Kaufmann.

VanLehn, K., Ball, W. & Kowalski, B. (1988). *Non-LIFO execution of cognitive procedures* (Tech. Rep. PCG-15). Dept. of Psychology, Carnegie-Mellon University.

Notes

[1]Some types of problems occur so often that their solution has become routine, and the subjects hardly notice that they have found and rectified a point of ignorance. For instance, students might not initially understand the referent of a mathematical symbol while reading a text or example, but after a few second's reflection, they retrieve (or construct) its meaning, and continue their reading. Presumably, they learn something from such an experience. The experience can be analyzed as a brief episode of problem solving, even though the subjects may not have thought of it as such.

[2]This proposal is similar to Anderson's Rational Analysis (Anderson, in press), except that the time scales and phenomena are different. Anderson seeks to explain the fixed, unchanging part of a person's mind -- the cognitive architecture -- by assuming that it is the product of genetic adaptation to the demands of the environment. The proposal here is to explain an individual's knowledge as the product of adaptation to the environment that has been experienced since birth.

[3]This theory is slightly more general than ACT* and SOAR. Those theories claim that people invariably select one of the unsatisfied goals that was created most recently (i.e., both ACT* and SOAR have a last-in-first-out goal stacks).

[4]CIRRUS does not understand natural language; the protocol must be encoded by humans before giving it to CIRRUS.

[5]Although SAPA was build to handle only thermodynamics, it could be redesigned to have more task generality by allowing the user to input a problem space.